

CLAIMS**What is claimed is:**

- 5 1. An apparatus for designing constraint-based systems, the apparatus comprising:
 logic configured to perform a system design algorithm, wherein when the first logic performs the system design algorithm, individual workloads are clustered to generate a cluster workload, and wherein clustering the individual workloads into a
 10 cluster workload enables the total number of constraint calculations that need to be performed to be reduced by performing constraint-based calculations on a per-cluster basis rather than on a per-workload basis.
2. An apparatus for designing storage systems, the apparatus comprising:
 15 logic configured to perform a storage system design algorithm, wherein when the first logic performs the storage system design algorithm, stores are clustered to generate a cluster workload, and wherein clustering stores into a cluster workload enables the total number of constraint calculations that need to be performed to be reduced.
- 20 3. The apparatus of claim 2, wherein the stores are clustered in such a way that the total number of stores is reduced.
4. The apparatus of claim 2, wherein the stores are clustered into aggregates
 25 having storage capacity requirements that prevent more than a particular number, N , of said aggregates to be placed per logical unit (LU).

5. The apparatus of claim 2, wherein the stores are clustered in such a way that the cluster workload comprises aggregates that have bandwidth requirements that prevent more than a particular number, N , of said aggregates to be placed per logical unit (LU).

5

6. The apparatus of claim 2, wherein the stores are clustered into aggregates having storage capacity requirements and bandwidth requirements that prevent more than a particular number, N , of said aggregates to be placed per logical unit (LU).

10

7. The apparatus of claim 2, wherein the logic is comprised by a computer and wherein the storage system design algorithm is implemented as a software program being executed by the computer.

15

8. The apparatus of claim 2, wherein the logic comprises a hardware configuration and wherein the storage system design algorithm is performed solely in hardware.

20

9. The apparatus of claim 2, wherein the stores are ordered in the cluster in the order in which they are presented to the logic.

10. The apparatus of claim 2, wherein the stores are ordered in the cluster randomly with respect to the order in which they are presented to the logic.

25

11. The apparatus of claim 2, wherein the stores are ordered in the cluster after the stores have been sorted based on characteristics of the stores.

12. The apparatus of claim 11, wherein the characteristics include size of the stores.
13. The apparatus of claim 11, wherein the characteristics include aggregate
5 stream demands.
14. The apparatus of claim 2, wherein the stores are ordered in the cluster by grouping together stores with similar attributes.
- 10 15. The apparatus of claim 2, wherein the stores are ordered in the cluster by taking into account desired properties of the cluster that is being created.
16. A method for designing storage systems, the method comprising the steps of:
clustering stores in accordance with a storage system design algorithm to produce a
15 cluster workload that reduces a number of constraint calculations that need to be performed per store:
ordering stores to be clustered;
clustering the stores;
assigning the cluster workload; and
20 de-clustering the workload into its component stores.
17. The method of claim 16, wherein the step of clustering includes the step of clustering stores together to reduce the total number of stores.
- 25 18. The method of claim 16, wherein the step of clustering includes the step of clustering stores into aggregates having storage capacity requirements that allow no more than a particular number of said aggregates to be placed per logical unit (LU).

19. The method of claim 16, wherein the step of clustering includes the step of clustering stores into aggregates having bandwidth requirements that allow no more than a particular number of said aggregates to be placed per logical unit (LU).

5

20. The method of claim 16, wherein the step of clustering includes the step of clustering the stores of the workload into aggregates having storage capacity requirements and bandwidth requirements that allow no more than a particular number of said aggregates to be placed per logical unit (LU).

21. The method of claim 16, wherein the step of ordering the stores is performed in one of the following ways:

(1) ordering the stores in the order in which they are presented to the design algorithm;

5 (2) randomly re-ordering the stores after the stores have been presented to the design algorithm;

(3) re-ordering the stores after the stores have been presented to the design algorithm by sorting the stores based on characteristics associated with the stores;

10 (4) re-ordering the stores after the stores have been presented to the design algorithm by grouping together stores with similar attributes; or

(5) re-ordering the stores after the stores have been presented to the design algorithm by taking into account desired properties of the cluster being created in the clustering step.

15 22. A computer program for designing storage systems, the computer program being embodied on a computer readable medium, the computer program comprising:

a first code segment for clustering stores to produce a cluster workload that reduces a number of constraint calculations that need to be performed per store;

a second code segment for assigning the cluster workload; and

20 a third code segment for de-clustering the workload into its component stores.